

Ref #	Hits	Search Query	DBs	Default Operator	Plurals	Time Stamp
L1	2	"20040107264"	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/12/07 14:19
L2	2	"20040107264".pn.	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/12/07 14:20
L3	159	(Yoshihiro near2 Nishida).in.	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/12/07 14:22
L4	1177	multiprocessor near2 environment	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/12/07 14:24
L5	25983	shared near2 memory	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/12/07 14:27
L6	66062	uninterrupt\$4	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/12/07 14:27
L7	638369	interrupt\$4	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/12/07 14:35
L8	26666	attribute\$3 near2 data	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/12/07 14:35
L9	759	(first adj processor) same (second adj processor) same share\$4	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/12/07 14:36
L10	27187	"711"/\$.ccls.	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/12/07 14:36

L11	310	9 and 10	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/12/07 14:36
L12	146	11 and 7	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/12/07 14:36
L13	130	5 and 12	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/12/07 14:36
L14	2	8 and 13	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/12/07 14:36


[Subscribe \(Full Service\)](#) [Register \(Limited Service, Free\)](#) [Login](#)

 Search: ☒ The ACM Digital Library ☐ The Guide


[Feedback](#) [Report a problem](#) [Satisfaction survey](#)

Terms used

[memory](#) [arbitration](#) [shared](#) [memory](#) [interrupt](#) [attribute](#) [data](#)

Found 113 of 167,655

Sort results by


[Save results to a Binder](#)
[Try an Advanced Search](#)
[Try this search in The ACM Guide](#)

Display results


[Search Tips](#)
☐ Open results in a new window

Results 1 - 20 of 113

Result page: [1](#) [2](#) [3](#) [4](#) [5](#) [6](#) [next](#)Relevance scale ☐ ☐ ☐ ☐ ☐

1 [Accelerating shared virtual memory via general-purpose network interface support](#)



Angelos Bilas, Dongming Jiang, Jaswinder Pal Singh

February 2001 **ACM Transactions on Computer Systems (TOCS)**, Volume 19 Issue 1

Publisher: ACM Press

Full text available: pdf(178.88 KB)

 Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#), [review](#)

Clusters of symmetric multiprocessors (SMPs) are important platforms for high-performance computing. With the success of hardware cache-coherent distributed shared memory (DSM), a lot of effort has also been made to support the coherent shared-address-space programming model in software on clusters. Much research has been done in fast communication on clusters and in protocols for supporting software shared memory across them. However, the performance of software virtual memory (SVM) is sti ...

Keywords: applications, clusters, shared virtual memory, system area networks

2 [VLSI assist for a multiprocessor](#)



Bob Beck, Bob Kasten, Shreekanth Thakkar

 October 1987 **ACM SIGARCH Computer Architecture News , ACM SIGPLAN Notices , ACM SIGOPS Operating Systems Review , Proceedings of the second international conference on Architectural support for programming languages and operating systems ASPLOS-II**, Volume 15 , 22 , 21 Issue 5 , 10 , 4

Publisher: IEEE Computer Society Press, ACM Press

Full text available: pdf(1.28 MB)

 Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Multiprocessors have long been of interest to computer community. They provide the potential for accelerating applications through parallelism and increased throughput for large multi-user system. Three factors have limited the commercial success of multiprocessor systems; entry cost, range of performance, and ease of application. Advances in very large scale integration (VLSI) and in computer aided design (CAD) have removed these limitations, making possible a new class of multiprocessor system ...

3 [Memory system performance of UNIX on CC-NUMA multiprocessors](#)



John Chapin, A. Herrod, Mendel Rosenblum, Anoop Gupta

 May 1995 **ACM SIGMETRICS Performance Evaluation Review , Proceedings of the 1995 ACM SIGMETRICS joint international conference on Measurement**

and modeling of computer systems SIGMETRICS '95/PERFORMANCE '95,
Volume 23 Issue 1

Publisher: ACM Press

Full text available:  [pdf\(1.78 MB\)](#)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

This study characterizes the performance of a variant of UNIX SVR4 on a large shared-memory multiprocessor and analyzes the effects of possible OS and architectural changes. We use a nonintrusive cache miss monitor to trace the execution of an OS-intensive multiprogrammed workload on the Stanford DASH, a 32-CPU CC-NUMA multiprocessor (CC-NUMA multiprocessors have cache-coherent shared memory that is physically distributed across the machine). We find that our version of UNIX accounts for 24% of ...

4 Experience Using Multiprocessor Systems—A Status Report



Anita K. Jones, Peter Schwarz

June 1980 **ACM Computing Surveys (CSUR)**, Volume 12 Issue 2

Publisher: ACM Press

Full text available:  [pdf\(4.48 MB\)](#)

Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#)

5 A survey of commercial parallel processors



Edward Gehringer, Janne Abullarade, Michael H. Gulyin

September 1988 **ACM SIGARCH Computer Architecture News**, Volume 16 Issue 4

Publisher: ACM Press

Full text available:  [pdf\(2.96 MB\)](#)

Additional Information: [full citation](#), [abstract](#), [citations](#), [index terms](#)

This paper compares eight commercial parallel processors along several dimensions. The processors include four shared-bus multiprocessors (the Encore Multimax, the Sequent Balance system, the Alliant FX series, and the ELXSI System 6400) and four network multiprocessors (the BBN Butterfly, the NCUBE, the Intel iPSC/2, and the FPS T Series). The paper contrasts the computers from the standpoint of interconnection structures, memory configurations, and interprocessor communication. Also, the share ...

6 Cache memory performance in a unix enviroment



Cedell Alexander, William Keshlear, Furrokh Cooper, Faye Briggs

June 1986 **ACM SIGARCH Computer Architecture News**, Volume 14 Issue 3

Publisher: ACM Press

Full text available:  [pdf\(2.10 MB\)](#)

Additional Information: [full citation](#), [citations](#), [index terms](#)

7 Multi-level shared caching techniques for scalability in VMP-M/C



D. R. Cheriton, H. A. Goosen, P. D. Boyle

April 1989 **ACM SIGARCH Computer Architecture News , Proceedings of the 16th**

annual international symposium on Computer architecture ISCA '89, Volume 17 Issue 3

Publisher: ACM Press

Full text available:  [pdf\(1.27 MB\)](#)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

The problem of building a scalable shared memory multiprocessor can be reduced to that of building a scalable memory hierarchy, assuming interprocessor communication is handled by the memory system. In this paper, we describe the VMP-MC design, a distributed parallel multi-computer based on the VMP multiprocessor design, that is intended to provide a set of building blocks for configuring machines from one to several thousand processors. VMP-MC uses a memory hierarchy based on shared caches ...

8 Memory systems: Cluster miss prediction with prefetch on miss for embedded CPU instruction caches

Ken Batcher, Robert Walker

September 2004 **Proceedings of the 2004 international conference on Compilers, architecture, and synthesis for embedded systems**

Publisher: ACM Press

Full text available:  pdf(343.66 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

Soft CPU cores are often used in embedded systems, yet they limit opportunities to improve cache performance to hardware assistance outside the CPU core. Instruction prefetching is commonly used, but the popular Prefetch On Miss (POM) technique is less helpful when the instruction flow does not follow a sequential execution order, which is often the case in real-time networking applications. Cluster Miss Prediction (CMP) can help in those worst case situations when cache misses do not follow a s ...

Keywords: WCET, cache design, cache prefetch, embedded systems, hiding memory latency, networking

9 Anatomy of a message in the Alewife multiprocessor

John Kubiawicz, Anant Agarwal

August 1993 **Proceedings of the 7th international conference on Supercomputing**

Publisher: ACM Press

Full text available:  pdf(1.36 MB) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)


Shared-memory provides a uniform and attractive mechanism for communication. For efficiency, it is often implemented with a layer of interpretive hardware on top of a message-passing communications network. This interpretive layer is responsible for data location, data movement, and cache coherence. It uses patterns of communication that benefit common programming styles, but which are only heuristics. This suggests that certain styles of communication may benefit from direct access to the ...

10 The minerva multi-microprocessor

Lawrence C. Widdoes

January 1976 **ACM SIGARCH Computer Architecture News , Proceedings of the 3rd annual symposium on Computer architecture ISCA '76**, Volume 4 Issue 4

Publisher: ACM Press

Full text available:  pdf(651.29 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)


A multiprocessor system is described which is an experiment in low cost, extensible, multiprocessor architectures. Global issues such as inclusion of a central bus, design of the bus arbiter, and methods of interrupt handling are considered. The system initially includes two processor types, based on microprocessors, and these are discussed. Methods for reducing processor demand for the central bus are described.

11 A characterization of sharing in parallel programs and its application to coherency protocol evaluation

S. J. Eggers, R. H. Katz

May 1988 **ACM SIGARCH Computer Architecture News , Proceedings of the 15th Annual International Symposium on Computer architecture ISCA '88**, Volume 16 Issue 2

Publisher: IEEE Computer Society Press, ACM Press

Full text available:  pdf(1.38 MB) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

In this paper we use trace-driven simulation to analyze the memory reference patterns of

write shared data in several parallel applications. We first develop a characterization of write sharing (based on the notion of a write run), and then examine the traces, using metrics derived from the characterization. The results indicate that the amount of write sharing in all programs is small; and that it is characterized by short to medium sequences of per processor references, with little conten ...

12 Synchronizing processors through memory requests in a tightly coupled



multiprocessor

A. Seznec, Y. Jégou

May 1988 **ACM SIGARCH Computer Architecture News , Proceedings of the 15th Annual International Symposium on Computer architecture ISCA '88**,
Volume 16 Issue 2

Publisher: IEEE Computer Society Press, ACM Press

Full text available: [pdf\(927.94 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

To satisfy the growing need for computing power, a high degree of parallelism will be necessary in future supercomputers. Up to the late 70s, supercomputers were either multiprocessors (SIMD-MIMD) or pipelined monoproducts. Current commercial products combine these two levels of parallelism. Effective performance will depend on the spectrum of algorithms which is actually run in parallel. In a previous paper [Je86], we have presented the DSPA processor, a pipeline processor whi ...

13 Computing curricula 2001



September 2001 **Journal on Educational Resources in Computing (JERIC)**

Publisher: ACM Press

Full text available: [pdf\(613.63 KB\)](#) [html\(2.78 KB\)](#) Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#)

14 The SHRIMP performance monitor: design and applications



Margaret Martonosi, Douglas W. Clark, Malena Mesarina

January 1996 **Proceedings of the SIGMETRICS symposium on Parallel and distributed tools**

Publisher: ACM Press

Full text available: [pdf\(1.01 MB\)](#) Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#)

15 The VMP multiprocessor: initial experience, refinements, and performance evaluation



D. R. Cheriton, A. Gupta, P. D. Boyle, H. A. Goosen

May 1988 **ACM SIGARCH Computer Architecture News , Proceedings of the 15th Annual International Symposium on Computer architecture ISCA '88**,
Volume 16 Issue 2

Publisher: IEEE Computer Society Press, ACM Press

Full text available: [pdf\(1.73 MB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

VMP is an experimental multiprocessor being developed at Stanford University, suitable for high-performance workstations and server machines. Its primary novelty lies in the use of software management of the per-processor caches and the design decisions in the cache and bus that make this approach feasible. The design and some uniprocessor trace-driven simulations indicating its performance have been reported previously. In this paper, we present our initial experience with the V ...

16 Hardware support for interprocess communication



 U. Ramachandran, M. Solomon, M. Vernon
June 1987 **Proceedings of the 14th annual international symposium on Computer architecture**

Publisher: ACM Press

Full text available:  [pdf\(1.10 MB\)](#)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

In recent years there has been increasing interest in message-based operating systems, particularly in distributed environments. Such systems consist of a small message-passing kernel supporting a collection of system server processes that provide such services as resource management, file service, and global communications. For such an architecture to be practical, it is essential that basic messages be fast, since they often replace what would be a simple procedure call o ...

17 Gilgamesh: a multithreaded processor-in-memory architecture for petaflops computing

Thomas L. Sterling, Hans P. Zima

November 2002 **Proceedings of the 2002 ACM/IEEE conference on Supercomputing**

Publisher: IEEE Computer Society Press

Full text available:  [pdf\(322.86 KB\)](#)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Processor-in-Memory (PIM) architectures avoid the von Neumann bottleneck in conventional machines by integrating high-density DRAM and CMOS logic on the same chip. Parallel systems based on this new technology are expected to provide higher scalability, adaptability, robustness, fault tolerance and lower power consumption than current MPPs or commodity clusters. In this paper we describe the design of *Gilgamesh*, a PIM-based massively parallel architecture, and elements of its execution mo ...

Keywords: Petaflops computing, Processor-In-Memory, data parallel processing, irregular applications, parallel architectures

18 Session summaries from the 17th symposium on operating systems principle (SOSP'99)

Jay Lepreau, Eric Eide

April 2000 **ACM SIGOPS Operating Systems Review**, Volume 34 Issue 2

Publisher: ACM Press

Full text available:  [pdf\(3.15 MB\)](#)

Additional Information: [full citation](#), [index terms](#)

19 Modeling and measurement of the impact of Input/Output on system performance

Janaki Akella, Daniel P. Siewiorek

April 1991 **ACM SIGARCH Computer Architecture News , Proceedings of the 18th annual international symposium on Computer architecture ISCA '91**, Volume 19 Issue 3

Publisher: ACM Press

Full text available:  [pdf\(952.88 KB\)](#)

Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#)

20 The VMP network adapter board (NAB): high-performance network communication for multiprocessors

H. Kanakia, D. Cheriton

August 1988 **ACM SIGCOMM Computer Communication Review , Symposium proceedings on Communications architectures and protocols SIGCOMM**

'88, Volume 18 Issue 4

Publisher: ACM Press

Full text available:  pdf(1.63 MB)Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

High performance computer communication between multiprocessor nodes requires significant improvements over conventional host-to-network adapters. Current host-to-network adapter interfaces impose excessive processing, system bus and interrupt overhead on a multiprocessor host. Current network adapters are either limited in function, wasting key host resources such as the system bus and the processors, or else intelligent but too slow, because of complex transport protocols and because of a ...

Results 1 - 20 of 113

Result page: [1](#) [2](#) [3](#) [4](#) [5](#) [6](#) [next](#)

The ACM Portal is published by the Association for Computing Machinery. Copyright © 2005 ACM, Inc.

[Terms of Usage](#) [Privacy Policy](#) [Code of Ethics](#) [Contact Us](#)Useful downloads:  [Adobe Acrobat](#)  [QuickTime](#)  [Windows Media Player](#)  [Real Player](#)